

An Entropy-based Approach to Wide Area Surveillance

Gaemus E. Collins^a, Mark R. Meloon^a, Kevin J. Sullivan^a, and Janice M. Chinn^b

^a Toyon Research Corp., 75 Aero Camino, Suite A, Goleta, CA, USA 93117;

^b AFRL/SNAR, Bldg 620, 2241 Avionics Circle, Wright-Patterson Air Force Base

ABSTRACT

The use of measures from information theory to evaluate the expected utility of a set of candidate actions is a popular method for performing sensor resource management. Shannon entropy is a standard metric for information. Past researchers have shown^{1–5} that the discrete entropy formula can measure the quality of identification information on a target, while the continuous entropy formula can measure kinematic state information of a target. In both cases, choosing controls to minimize an objective function proportional to entropy will improve ones information about the target. However, minimizing entropy does not naturally promote *detection* of new targets or “wide area surveillance” (WAS). This paper outlines a way to use Shannon entropy to motivate sensors to track (partially) discovered targets and survey the search space to discover new targets simultaneously. Results from the algorithmic implementation of this method show WAS being favored when most targets in the search space are undiscovered, and tracking of discovered targets being favored when most targets are in track. The tradeoff between these two competing objectives is adjusted by the objective function automatically and dynamically.

Keywords: Sensor Management, Entropy, Information Theory, Wide Area Surveillance

1. INTRODUCTION

Sensor resource management is the control mechanism in information fusion.⁶ Sensor schedules drive a system of sensors and processing algorithms to improve knowledge of a battlespace. Sensor management algorithms have been researched for many years. A number of techniques exist for solving the problem.^{5,7} One popular and successful approach has been the application of measures from information theory. These approaches typically entail calculating the amount of expected information resulting from a candidate sensor action to determine whether it merits execution. Hintz and McVey⁸ were the first to use information theory for sensor management. In that work, they use a measure of entropy due to Shannon⁴ to track multiple targets with a single sensor. The discrete entropy formula can measure the quality of identification (ID) information on a target, while the continuous entropy formula can measure the kinematic state information. In both cases, choosing controls to minimize an objective function proportional to entropy will improve the quality of information about the target.

While information-theoretic approaches have been shown to perform favorably in a variety of sensor management and target tracking problems, many of these problems have included only one goal. In a typical problem, however, it may be desirable to direct sensors to accomplish multiple goals simultaneously. For example, one may wish to both improve ID and kinematic information of a target at the same time, such as is done in.^{9,10} Directing sensors to improve information about currently tracked targets and detect new targets is another pair of complimentary goals. Achieving a balance between these two is especially important in difficult tracking situations such as long-duration missions, environments with significant obscuration, vehicles that undergo stop-move-stop motion, and lack of sensor availability. In these cases, it is unlikely that track can be maintained on all targets of importance during an entire mission. Sensor resources dedicated to improving, say, the ID of a currently tracked target are less resources available for performing “wide area surveillance” (WAS) to detect new targets or targets that have been lost. Current entropy-based approaches do not naturally handle this tradeoff between competing mission goals.

This paper outlines a way to use Shannon entropy to motivate sensors to track discovered targets and survey the search space to discover new targets simultaneously. In this method, an objective function is defined as a

Further author information, send correspondence to gcollins@toyon.com

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE APR 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE An Entropy-based Approach to Wide Area Surveillance				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Toyon Research Corporation, 75 Aero Camino Drive Suite A, Goleta, CA, 93117-3139				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Proceedings of the SPIE Conference on Signal Processing, Sensor Fusion, and Target Recognition XV (Orlando, FL), vol. 6235, April 2006. U.S. Government or Federal Rights License					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 13	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

weighted sum over the expected entropy values of all targets. A heterogeneous distribution of targets is assumed to exist in a spatially discretized search space. These undiscovered targets are added into the total objective function with maximum entropy, H_{\max} . Once detected, these targets will be added into the objective function with some value less than H_{\max} , thereby lowering the value of the objective function. To evaluate candidate controls, the expected value of the objective function (factoring in entropy from undiscovered targets and targets in track) is computed, and controls promising the most significant decrease to the objective function are chosen.

The paper begins with a discussion of the objective function used for evaluation of candidate sensor controls in section 2. The objective function presented here has two components, but only the WAS-based component will be discussed in detail. In section 3 a WAS-based objective function is derived to evaluate candidate sensor tasks under the assumption that they are independent. Then in section 4 this objective function is extended to evaluate entire schedules of tasks, including tasks that may be occurring simultaneously or with partial time overlap.

2. THE OBJECTIVE FUNCTION FOR SENSOR CONTROLS

In this paper we consider a general sensor scheduling problem in which tasks for multiple sensors at multiple times must be automatically queued to (a) maintain track information (identification and/or kinematics) on discovered targets, *and* (b) discover undiscovered targets in the search space. To this end, we employ an “information metric” or “objective function” to score possible candidate sensor tasks and chose the best ones. This metric is based on discrete Shannon entropy,⁴ and has two components corresponding to the two goals stated above:

- a track-based objective function, $J_{\mathcal{T}}$, that measures the quality of information on existing “tracks” (discovered targets), and
- a WAS-based objective function, $J_{\mathcal{W}}$, that measures the quality of information gained from WAS.

The two objective functions work together and naturally counterbalance each other to choose tasks for WAS when there are very few targets in track, and chose tasks for track maintenance when we believe most of the targets in the search space are discovered and in track. Overall, we seek a sensor schedule that will minimize the sum of the two competing terms

$$J = J_{\mathcal{T}} + J_{\mathcal{W}}.$$

The discrete Shannon entropy formula is a means of measuring the quality of an identification estimate. (While this paper works with the discrete entropy formula, the theory presented is compatible with the continuous form of entropy (for kinematics) as well. The formulas, however, require some modification.) Let $\mathbf{p} = [p_1, \dots, p_N]$ be a classification vector, where p_i is the probability that the target belongs to class i , $i \in \{1, \dots, N\}$. The entropy of p is defined as

$$H(\mathbf{p}) = - \sum_{i=1}^N p_i \cdot \log(p_i).$$

It is immediate that

- $H(\mathbf{p}) \geq 0$,
- $H_{\max} = \log(N)$, resulting from a completely unclassified target with probability vector $\mathbf{p} = [\frac{1}{N}, \dots, \frac{1}{N}]$,
- $H_{\min} = 0$, resulting from a perfectly classified target, with probability vector $\mathbf{p} = [0, \dots, 0, 1, 0, \dots, 0]$,
- better target classification corresponds to lower entropy.

For targets being tracked, each has a classification vector computed by accumulating evidence from measurements that are associated with a target through a fusion and tracking process. Targets assumed to exist in the search space undiscovered are given classification vectors $\mathbf{p} = [\frac{1}{N}, \dots, \frac{1}{N}]$, since no classification information exists for these targets.

In the objective function J , the entropy of each target is weighted by the target classification vector p in such a way that targets that are classified as “targets of interest” (TOIs) will affect the objective function more than

targets that are not classified as TOIs. Mathematically, this is given by $W(\mathbf{p}) \cdot H(\mathbf{p})$ where

$$W = W(\mathbf{p}) = \sum_{i \in \{\text{TOIs}\}} p_i.$$

The schedule evaluation process cannot have access to the results of the fusion and tracking process, as that would require the candidate tasks to be executed before they were evaluated. So the objective function is computed from *expected entropy* rather than actual entropy values. The expected entropy is a function of new classification information expected to be obtained from a candidate sensor task. We estimate the expected entropy by considering all of the possible classification outcomes weighted by their probability of occurring. For the l^{th} target, this produces

$$\hat{H}(\mathbf{p}^l) = \sum_{n \in \text{Outcome}} \hat{H}_{l|n} \hat{P}_l(n),$$

where “Outcome” is a declaration by a target classification process.

Targets that are believed to be discovered or partially known are called *tracks*, and these tracks have classification vectors \mathbf{p}^l that are maintained in a tracking module. The score of a candidate task is partially based on its ability to improve the classification of tracks in the tracking module. This is reflected in the track component of the score,

$$J_T = \sum_{\mathbf{p}^l \in \{\text{existing tracks}\}} W(\mathbf{p}^l) \hat{H}(\mathbf{p}^l).$$

Extensive work supports the use of entropy to measure target classification in this way¹⁻⁵ and therefore analysis of J_T will be excluded from this paper.

3. EVALUATION OF A SENSOR TASK

The motivation to perform wide area surveillance (WAS) to initiate new tracks stems from a belief that there exist additional (undiscovered) targets scattered throughout the search space that are not currently in track. The objective function should reward sensors for discovering these targets and tracking them. This assertion is formalized by hypothesizing a distribution of undiscovered targets and designing the objective function J_W that decreases as undiscovered targets are discovered.

Suppose $T(t_{\text{start}}, t_{\text{end}}, S)$ represents a single observation or task performed by sensor S from t_{start} to t_{end} . Further, suppose that the search space is divided into non-overlapping “gridblocks” B_j , and the footprint of task $T(t_{\text{start}}, t_{\text{end}}, S)$ covers some of the gridblocks. To avoid complicated polygon intersection calculations, a sensor footprint covering the center point of a gridblock (the “gridpoint”) will be modeled as coverage of that entire gridblock. This approximation will produce a computation error proportional to the area of the individual gridblocks, and this can be decreased in the simulation by dividing the search space into smaller gridblocks.

If the footprint of task $T(t_{\text{start}}, t_{\text{end}}, S)$ intersects the gridblock B_j , one expects that $T(t_{\text{start}}, t_{\text{end}}, S)$ should reveal some of the undiscovered targets in B_j . Tasks that reveal more undiscovered targets and/or tasks that gather more information about the discovered targets should be favored by J_W by lowering its value.

In what follows, the parameter “ T ” appearing with a quantity will designate the value of that quantity assuming T was executed; t_0 will be a time before the start time of task T , t_T will be the duration of the task T , t_1 will be a time after the completion time of task T , and $t_{\text{DIFF}} = t_1 - t_0$. Let:

- $N_{ij}^U(t, T)$, ($N_{ij}^U(t)$), be a random variable representing the number of undetected targets of type i located within gridblock B_j at time t with (without) the observation from task T ;
- $N_j^U(t, T) = \sum_i N_{ij}^U(t, T)$, ($N_j^U(t) = \sum_i N_{ij}^U(t)$);
- $\lambda_{ij}^U(t, T) = E[N_{ij}^U(t, T)]$, ($\lambda_{ij}^U(t) = E[N_{ij}^U(t)]$) be the expected value of $N_{ij}^U(t, T)$, ($N_{ij}^U(t)$);
- $\lambda_j^U(t, T) = E[N_j^U(t, T)]$, ($\lambda_j^U(t, \sim T) = E[N_j^U(t, \sim T)]$);

- $D_i(X_j, t, T)$ be the probability that if a target of type i does indeed exist at location X_j (within gridblock B_j) at time t that task T will detect it (obviously $D_i(X_j, t) = 0$ if X_j is not in the footprint of task T);
- $N_{ij}^D(t, T, N_{ij}^U(t))$ be a random variable representing the number of targets of type i within gridblock B_j that will be detected by task T at time t ;
- $E[N_{ij}^D(t, T, N_{ij}^U(t))]$ be the expected number of targets of type i within gridblock B_j that will be detected by task T at time t .

The mean $\lambda_{ij}^U(t)$ represents the sensor managers best estimate of how many targets of type i are in B_j , and can be formed on the basis of *a priori* information, data gathered during the course of the mission, or a combination of the two (e.g., *a priori* estimates refined using analysis of tracker data). In this paper, $\lambda_{ij}^U(t)$ will be estimated at t_0 and then maintained using the equations below.

To determine the optimal sensor task T at time t_0 , one would minimize

$$J(t_1, T) = J_{\mathcal{T}}(t_1, T) + J_{\mathcal{W}}(t_1, T),$$

where the dependence of T on t_{start} , t_{end} , and S has been suppressed. The weighted expected entropy of existing tracks (with the expected improvement from task T) is

$$J_{\mathcal{T}}(t_1, T) = \sum_{\mathbf{p}^l \in \{\text{existing tracks}\}} W(\mathbf{p}^l) \cdot \hat{H}(\mathbf{p}^l, t_1, T)$$

with $\hat{H}(\mathbf{p}^l, t_1, T)$ representing the expected entropy of the track \mathbf{p}^l at time t_1 incorporating task T . However, the weighted entropy of undiscovered and newly discovered targets represented in the $J_{\mathcal{W}}(t_1, T)$ term is the focus of this paper.

$J_{\mathcal{W}}(t_1, T)$ is a value of the objective function in the future, and is unavailable to the sensor manager. To evaluate *candidate* tasks, the sensor manager must estimate this value. This involves computing $\lambda_{ij}^U(t_1, T)$ from $\lambda_{ij}^U(t_0)$ for each possible candidate task T . Assuming that a sensor task T with footprint covering gridblock B_j should reveal some of the undiscovered targets in B_j , the value λ_{ij}^U in those gridblocks should be decreased, and those discovered targets should be transferred into new tracks. Meanwhile, gridblocks B_j not covered by the footprint of T may be accumulating undiscovered targets; thus, λ_{ij}^U may increase for these unobserved gridblocks as time goes by. These statements are formalized in the equation

$$\begin{aligned} J_{\mathcal{W}}(t_1, T) = & \sum_{j \notin \text{footprint of } T} W_j \cdot \lambda_{ij}^U(t_1, T) \cdot H_{\max} + \sum_{j \in \text{footprint of } T} W_j \cdot \lambda_{ij}^U(t_1, T) \cdot H_{\max} \\ & + \sum_{\mathbf{p}^l \in \{\text{new tracks}\}} W(\mathbf{p}^l) \cdot H(\mathbf{p}^l, t_1, T) \end{aligned}$$

with $\lambda_{ij}^U(t_1, T) = \lambda_{ij}^U(t_0) + \epsilon_{ij}(t_{\text{DIFF}})$ for $j \notin \text{footprint of } T$, and $\lambda_{ij}^U(t_1, T) = \lambda_{ij}^U(t_0) - K_{ij} + \epsilon_{ij}(t_{\text{DIFF}} - t_T)$ for $j \in \text{footprint of } T$. K_{ij} represents the number of new tracks of type i discovered by T in gridblock B_j , and the sum over new tracks runs from 1 to $\sum_{i,j} K_{ij}$. The factor $\epsilon_{ij}(s)$ is added to the undiscovered targets (hereafter referred to as “target mass” because the value may be fractional) while gridblocks are not being observed, to represent undiscovered targets of type i that might be accumulating in gridblock B_j while no sensor task is observing it. $\epsilon_{ij}(s)$ should be computed so that the undiscovered target mass in gridblock B_j would have bounded growth. For example, $\epsilon_{ij}(s) = (1 - 2^{-s})(\lambda_{ij}^U(0) - \lambda_{ij}^U(t_0))$ has the properties that $\epsilon_{ij}(0) = 0$ and $\lim_{s \rightarrow \infty} \epsilon_{ij}(s) = \lambda_{ij}^U(0) - \lambda_{ij}^U(t_0)$; so the undiscovered target mass in an unobserved cell will eventually grow to its original value at the beginning of the simulation.

The new tracks, \mathbf{p}^l , will not actually be *known* prior to the execution of task T ; therefore, values dependent on these new tracks must be estimated by the sensor manager. The number of new tracks of type i originating from gridblock B_j as a result of task T is estimated as

$$K_{ij} \approx E[N_{ij}^D(t, T, N_{ij}^U(t))].$$

Therefore $\lambda_{ij}^U(t_1, T)$, the average number of undiscovered targets following task T , must be approximated in gridblocks covered by T as

$$\lambda_{ij}^U(t_1, T) \approx \lambda_{ij}^U(t_0) - E[N_{ij}^D(t, T, N_{ij}^U(t))] + \epsilon_j(t_{\text{DIFF}} - t_T)$$

for $j \in \text{footprint of } T$. The entropy of the new tracks originating from gridblock B_j as a result of task T will be represented by an expected entropy,

$$H(p^l, t_1, T) \approx \hat{H}_j(t_1, T).$$

This value can be approximated based on the type of sensor being used, the geometry of the task T , etc. The sensor manager's estimate of the entropy from newly discovered tracks (by task T) becomes

$$\sum_{\mathbf{p}^l \in \{\text{new tracks}\}} W(\mathbf{p}^l) \cdot H(\mathbf{p}^l, t_1, T) \approx \sum_{j \in \text{footprint of } T} \left\{ W_j \cdot E[N_{ij}^D(t, T, N_{ij}^U(t))] \cdot \hat{H}(\mathbf{p}^l, t_1, T) \right\}.$$

Based on these approximations, The complete objective function can be rewritten as

$$\begin{aligned} J(t_1, T) \approx & \sum_{\mathbf{p}^l \in \{\text{existing tracks}\}} W(\mathbf{p}^l) \cdot \hat{H}(\mathbf{p}^l, t_1, T) + \sum_{j \notin \text{footprint of } T} W_j \cdot (\lambda_j^U(t_0) + \epsilon_j(t_{\text{DIFF}})) \cdot H_{\max} \\ & + \sum_{j \in \text{footprint of } T} \left\{ W_j \cdot (\lambda_j^U(t_0, T) - E[N_{ij}^D(t, T, N_{ij}^U(t))] + \epsilon_{ij}(t_{\text{DIFF}} - t_T)) \cdot H_{\max} \right\} \\ & + \sum_{j \in \text{footprint of } T} \left\{ W_j \cdot (E[N_{ij}^D(t, T, N_{ij}^U(t))] \cdot \hat{H}(\mathbf{p}^l, t_1, T)) \right\}, \end{aligned} \quad (1)$$

where “existing tracks” refer to tracks that existed prior to/independent of task T , and $\epsilon_j = \sum_i \epsilon_{ij}(s)$.

The expected value $E[N_{ij}^D(t, T, N_{ij}^U(t))]$ is computed from the law of total expectation. If $E[N_{ij}^D(t, T, N_{ij}^U(t)) | N_{ij}^U(t) = n]$ represents the expected number of targets within gridblock B_j that will be detected by task T at time t , given that exactly n targets are present, then

$$\begin{aligned} E[N_{ij}^D(t, T, N_{ij}^U(t))] &= \sum_{n=0}^{\infty} \{ E[N_{ij}^D(t, T, N_{ij}^U(t)) | N_{ij}^U(t) = n] \cdot P(N_{ij}^U(t) = n) \} \\ &= \sum_{n=0}^{\infty} \sum_{k=0}^n \{ P(N_{ij}^D(t, T, N_{ij}^U(t)) = k | N_{ij}^U(t) = n) \cdot k \cdot P(N_{ij}^U(t) = n) \}. \end{aligned}$$

From here on it is convenient, and often reasonable, to assume that

- detections occur independent of one another,
- detection probability for each target type is constant throughout a gridblock,
- detection probability for each target type is constant in time (within each scheduling interval).

Under these assumptions, $D_i(X_j, t, T(t, S)) = D_{ij}(T(S))$, and detections can be modeled by a binomial distribution, so that

$$P(N_{ij}^D(t, T, N_{ij}^U(t)) = k | N_{ij}^U(t) = n) = \binom{n}{k} \cdot (D_{ij}(T(S)))^k \cdot (1 - D_{ij}(T(S)))^{n-k}.$$

In this paper it will also be assumed that

- the probability of detection will not depend on target type, i.e. $D_{ij}(T(S)) = D_j(T(S))$.

(This is not necessary for the theory, but makes the formulas much more readable.)

In the absence of additional information about the distribution of undiscovered targets within a gridblock, it is practical to assume that N_{ij}^U follows a Poisson distribution, so that

$$P(N_{ij}^U(t) = n) = \frac{(\lambda_{ij}^U(t))^n \cdot \exp(-\lambda_{ij}^U(t))}{n!}.$$

Under these assumptions the formula for the expected number of targets of type i discovered gridblock B_j becomes

$$\begin{aligned} E[N_{ij}^D(t, T, N_{ij}^U(t))] &= \sum_{n=0}^{\infty} \sum_{k=0}^n \left\{ \binom{n}{k} \cdot (D_j(T(S)))^k \cdot (1 - D_j(T(S)))^{n-k} \cdot \left(\frac{k}{n!}\right) \cdot (\lambda_{ij}^U(t))^n \cdot \exp(-\lambda_{ij}^U(t)) \right\} \\ &= D_j(T(S)) \cdot \lambda_{ij}^U(t). \end{aligned}$$

The objective function becomes

$$\begin{aligned} J(t_1, T) &= \sum_{\mathbf{p}^l \in \text{existing tracks}} W(\mathbf{p}^l) \cdot \hat{H}(\mathbf{p}^l, t_1, T) + \sum_{j \notin \text{footprint of } T} W_j \cdot (\lambda_j^U(t_0) + \epsilon_j(t_{\text{DIFF}})) \cdot H_{\max} \\ &+ \sum_{j \in \text{footprint of } T} \left\{ W_j \cdot [(1 - D_j(T)) \cdot \lambda_j^U(t_0, T) + \epsilon_j(t_{\text{DIFF}} - t_T)] \cdot H_{\max} + D_j(T) \cdot \lambda_j^U(t_0, T) \cdot \hat{H}_j(t_1, T) \right\}. \end{aligned} \quad (2)$$

The term $W_j \cdot ((1 - D_j(T)) \cdot \lambda_j^U(t_0, T) + \epsilon_j(t_{\text{DIFF}} - t_T)) \cdot H_{\max}$ in (2) represents weighted entropy of targets within the footprint of task T that escape discovery by T . The term $W_j (D_j(T) \cdot \lambda_j^U(t_0, T) \cdot \hat{H}_j(t_1, T))$ in (2) represents the (unweighted) entropy of targets newly discovered by task T . This function should be minimized to find the best candidate tasks T between t_0 and t_1 .

From the objective function, it is easy to identify how the components of a task contribute to its score. The first sum, over pre-existing tracks, is made small by choosing tasks that improve the sensor managers ID information on existing tracks. The second sum, over gridblocks not in the observation footprint, is positive and grows with time. This quantity is made small by choosing tasks with a large footprint and/or long duration. The final sum, over gridblocks within the observation footprint, is made small by choosing tasks that cover gridblocks with many undetected targets and/or gridblocks with a high probability of detecting undiscovered targets. This score formula also displays the balance between searching with small footprint sensors that provide better classification information (e.g. EO/IR camera) and searching with large footprint sensors that cover more gridblocks (e.g. GMTI radar).

4. EVALUATION OF A SENSOR SCHEDULE

A sensor schedule is a list of sensor tasks for each sensor, spanning a common time interval. In this section, $[t_0, t_1]$ will be a scheduling interval. t_1 is referred to as the “horizon time” of the schedule. Equation (2) was derived to score a single sensor task under the implicit assumption that there were no other tasks occurring between t_0 and t_1 . In deriving a formula to score a sensor schedule over $[t_0, t_1]$, it is no longer possible to assume that tasks are completely isolated and occurring independent from one another. Equation (2) must be generalized.

Suppose schedule Ψ consists of tasks $\{T_1, \dots, T_M\}$ to be executed by one or more sensors. If Ψ represents a multi-sensor schedule, then some of the T_1, \dots, T_M sensor tasks may be occurring simultaneously. Simulating and scoring simultaneous tasks presents a difficult hurdle for many information-based tracking methods. A significant advantage of the method presented in this paper is that the objective function gracefully extends to scoring multi-sensor, multi-task sensor schedules, even those including simultaneous or partially time-overlapping tasks.

A multi-task objective function should be useful in different types of sensor scheduling algorithms. Scheduling algorithms often fall into two categories: “exhaustive” and “greedy” (i.e. branch-and-bound). The exhaustive scheduling algorithm considers all possible sensor schedules (combinations of tasks) over a particular scheduling time interval $[t_0, t_1]$. Greedy algorithms limit the set of candidate schedules considered by making a final task decision at each decision point.

Figure 1(a) is a tree diagram displaying all the candidate sensor schedules that an exhaustive scheduling algorithm might examine. At the first decision point, the sensor manager must choose between competing candidate tasks T_{11} , T_{12} and T_{13} . At the second decision point, the sensor manager must choose between

competing candidate tasks T_{21}, \dots, T_{27} (depending on which of T_{11}, \dots, T_{13} is being considered). The complete candidate schedules are 2-task sequences of the form $\{T_{1i}, T_{2j}\}$, $i = 1, 2, 3, j = 1, \dots, 7$, and are shown as leaf nodes in the tree (far right). The exhaustive scheduling algorithm generates each complete schedule, then evaluates the schedule with an objective function.

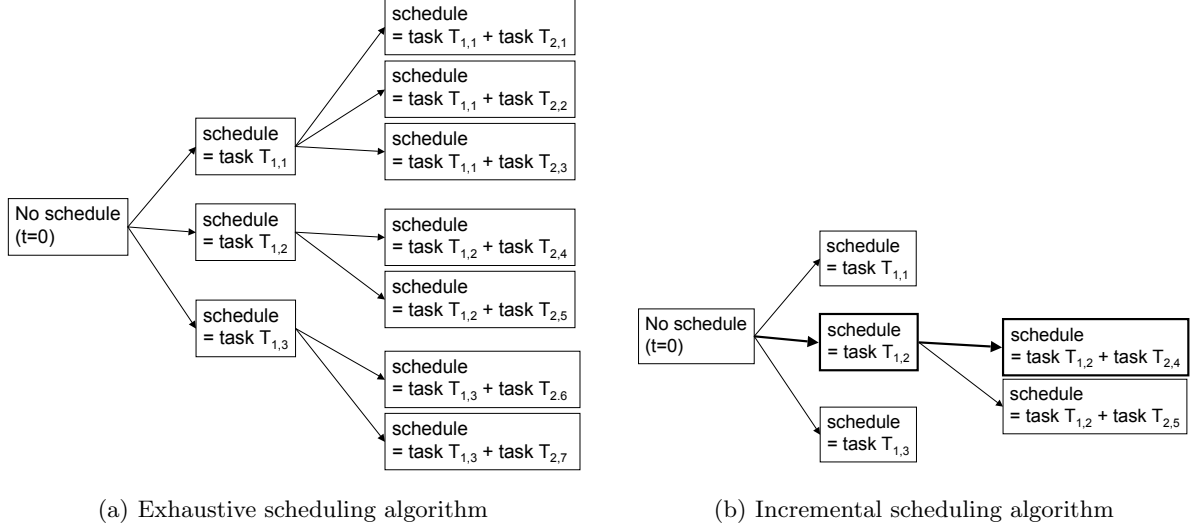


Figure 1. Search trees for exhaustive and incremental scheduling algorithms

In Figure 1(b), the same scenario is searched with a greedy algorithm. Tasks T_{11} , T_{12} , and T_{13} are individually evaluated with an objective function, and T_{12} is selected for the sensor schedule (as shown in **bold**). Then tasks T_{24} and T_{25} are individually evaluated, and T_{24} is selected. Schedules including tasks T_{21} , T_{22} , T_{23} , T_{26} , and T_{27} are never considered, since task T_{12} was selected at the first decision point.

The extension of (2) to an objective function for multi-task sensor schedules depends on the type of scheduling algorithm employed. We first consider the exhaustive case. The multi-task track-based objective function must reflect the potential improvement to every track estimate by every task, although only certain tasks will cover (and therefore improve the state estimate of) certain tracks. The track-based portion of the multi-task objective function is

$$J_T = \sum_{\mathbf{p}^l \in \{\text{existing tracks}\}} W(\mathbf{p}^l) \cdot \hat{H}(\mathbf{p}^l, t_1, T_1, \dots, T_M).$$

To describe a WAS-based objective function for the exhaustively generated multi-task schedule, some extra notation will be required. Let

$$d_j^{k_T}(T) = \begin{cases} (1 - D_j(T)) & \text{for } k_T = 0 \\ D_j(T) & \text{for } k_T = 1 \end{cases}, \quad \hat{H}_j(t, k_T) = \begin{cases} \hat{H}_j(t) & \text{for } k_T = 0 \\ \hat{H}_j(t, T) & \text{for } k_T = 1 \end{cases}. \quad (3)$$

The second and third terms of (2) represent entropy of target mass that remains in the search space undiscovered by T , and entropy of targets expected to be discovered by T . With multiple tasks in a schedule, undiscovered target mass in the search space may be discovered by up to M tasks. Target mass that remains undiscovered at t_1 (following the evaluation of all tasks in the schedule) must avoid detection by all tasks in the schedule. The multi-task WAS-based objective function uses the notation of (3) to represent all these possibilities. The sum over expected entropy of new detections or non-detections has 2^m sums, one for each combination of discovered/not discovered by task T_m , $m \in \{1, \dots, M\}$:

$$\sum_j \left\{ W_j \cdot \lambda_j^U(t_0) \cdot \sum_{k_{T_1}, \dots, k_{T_M}=0}^1 \left\{ d_j^{k_{T_1}}(T_1) \cdots d_j^{k_{T_M}}(T_M) \cdot \hat{H}_j(t_1, k_{T_1}, \dots, k_{T_M}) \right\} \right\}.$$

The complete objective function for the exhaustively generated multi-task sensor schedule $\{T_1, \dots, T_M\}$ is therefore given by

$$J = \sum_{\mathbf{p}^l \in \{\text{existing tracks}\}} W(\mathbf{p}^l) \cdot \hat{H}(\mathbf{p}^l, t_1, T_1, \dots, T_M) \quad (4)$$

$$+ \sum_j \left\{ W_j \cdot \lambda_j^U(t_0) \cdot \sum_{k_{T_1}, \dots, k_{T_M}=0}^1 \left\{ d_j^{k_{T_1}}(T_1) \cdots d_j^{k_{T_M}}(T_M) \cdot \hat{H}_j(t_1, k_{T_1}, \dots, k_{T_M}) \right\} \right\} \\ + \sum_j \{W_j \cdot \epsilon_j(\tau_j) \cdot H_{\max}\}. \quad (5)$$

The multi-task objective function associated with greedy scheduling algorithms should be designed to score individual tasks that are potentially occurring simultaneously with other tasks. A schedule consisting of tasks $\{T_1, \dots, T_M\}$ will be generated incrementally, and each task in the schedule must be compared to a set of competing candidate tasks at each particular decision point. Suppose that tasks $\{T_1, \dots, T_{K-1}\}$ have been selected by the greedy scheduling algorithm. We seek an objective function with which to score candidate tasks T_K that would follow T_{K-1} in the schedule.

The track-based objective function is the same as (2) except that it must account for the effects of tasks $\{T_1, \dots, T_{K-1}\}$ on the tracks,

$$J_{\mathcal{T}} = \sum_{\mathbf{p}^l \in \{\text{existing tracks}\}} W(\mathbf{p}^l) \cdot \hat{H}(\mathbf{p}^l, t_1, T_1, \dots, T_{K-1}, T_K).$$

The WAS-based objective function accounts for tasks $\{T_1, \dots, T_{K-1}\}$ by removing the expected discoveries by those tasks from the target mass distribution. Candidate task T_K has zero probability of detecting targets that were expected to be detected by tasks $\{T_1, \dots, T_{K-1}\}$. The expected discoveries of T_K are based on the target mass that remains in the search space undiscovered by $\{T_1, \dots, T_{K-1}\}$,

$$\lambda_j^U(0) \cdot (1 - D_j(T_1)) \cdots (1 - D_j(T_{K-1})).$$

The weighted entropy of targets expected to be discovered by T_K is therefore

$$\sum_j \left\{ W_j \cdot \lambda_j^U(0) \cdot (1 - D_j(T_1)) \cdots (1 - D_j(T_{K-1})) \cdot D_j(T_K) \cdot \hat{H}_j(t_1, T_K) \right\}.$$

The expected weighted entropy of targets remaining in the search space undiscovered by T_K is given by

$$\sum_j \left\{ W_j \cdot \lambda_j^U(0) \cdot (1 - D_j(T_1)) \cdots (1 - D_j(T_K)) \cdot H_{\max} \right\}.$$

Therefore, to compare competing candidate tasks at decision point T_K , it is necessary to compute (for each candidate task T_K)

$$J(t_1, T_1, \dots, T_K) \approx \sum_{\mathbf{p}^l \in \{\text{existing tracks}\}} \left\{ W(\mathbf{p}^l) \cdot \hat{H}(\mathbf{p}^l, t_1, T_1, \dots, T_K) \right\} \\ = \sum_j \left\{ W_j \cdot \lambda_j^U(0) \cdot \left(\prod_{k=1}^{K-1} (1 - D_j(T_k)) \right) \cdot D_j(T_K) \cdot \hat{H}_j(t_1, T_K) \right\} \\ + \sum_j \left\{ W_j \cdot \lambda_j^U(0) \cdot \left(\prod_{k=1}^K (1 - D_j(T_k)) \right) \cdot H_{\max} \right\} + \sum_j \{W_j \cdot \epsilon_j(\tau_j) \cdot H_{\max}\}.$$

The objective function for the entire schedule $\{T_1, \dots, T_M\}$ includes (a) the weighted entropy of every track following the evaluation of $\{T_1, \dots, T_M\}$, (b) $(m+1)$ sums over expected entropy of new detections, one for the expected detections of each task, and one for the targets not expected to be discovered by any task, and (c) and undiscovered target mass growth term. The incremental score of schedule Ψ is therefore

$$\begin{aligned} J(t_1, T_1, \dots, T_M) &\approx \sum_{\mathbf{p}^l \in \{\text{existing tracks}\}} \left\{ W(\mathbf{p}^l) \cdot \hat{H}(\mathbf{p}^l, t_1, T_1, \dots, T_M) \right\} \\ &= \sum_j \left\{ W_j \cdot \lambda_j^U(0) \cdot \left[\sum_{m=1}^M \left(\prod_{K=1}^{m-1} (1 - D_j(T_K)) \right) \cdot D_j(T_m) \cdot \hat{H}_j(t_1, T_m) \right] \right\} \\ &+ \sum_j \left\{ W_j \cdot \lambda_j^U(0) \cdot (1 - D_j(T_1)) \cdots (1 - D_j(T_M)) \cdot H_{\max} \right\} + \sum_j \left\{ W_j \cdot \epsilon_j(\tau_j) \cdot H_{\max} \right\}. \end{aligned}$$

5. SIMULATION RESULTS




Figure 2. SLAMEM™ screen capture of scenario region and airborne radar platforms

To demonstrate the feasibility of the sensor management algorithm, the logic was integrated into Toyon's SLAMEM™ simulation and used to control the sensors in an example military surveillance mission. The mission is to maintain continuous track of several targets of interest (TOIs) as they move about an area of interest (AOI). The AOI is shown in Figures 3(a) and 3(b). The images are screen captures of a SLAMEM™ display showing the terrain and road network (note the scale in the bottom right of the display). The AOI chosen for this simulation features complex terrain, with a mix of hills and level plains. Tracking of vehicles in hills and mountainous regions is difficult because the sensors may not have clear line of sight (LOS) to the targets.

A measure of effectiveness (MOE) of the tracking and surveillance algorithms is the number of TOIs that the automatic tracker has in track that are correctly classified (defined as the probability of being a tracked vehicle greater than 0.6). Although this is a tracking metric, the performance of the tracker relies heavily on the sensor measurements that are obtained. Because our sensor resource management algorithm is designed to provide the necessary sensor coverage to support long-term track identification and maintenance of targets, this tracking metric is a natural MOE to focus on. In the plots that follow, the curves are the averaged results of ten Monte Carlo trials in SLAMEM™.

All test cases include two airborne assets flying in orbits (one of which is shown as a purple line in Figure 2) approximately 150km away from the center of the AOI, and phased so that the sensor of one asset is always available whenever the other asset is in a turn (the sensors are set up to be untaskable when the asset is turning). Each asset contains a multi-mode radar; some radar modes are capable of providing classification information which is used as evidence in a Bayesian ID scheme, whereas other modes only provide kinematic data for use in updating the track.

The algorithm considers candidate radar tasks 10 degrees in azimuth. A new sensor schedule is developed every 15 seconds. Undiscovered target densities for the calculation of $J_{\mathcal{W}}$ were set between 0.01 and 0.7 per

square kilometer, depending on target type. The detectability $D_i(T)$ was fixed at $(0.9 * 0.85) = 0.765$, assuming a P_D of 0.9 and a probability of undiscovered target in motion of 0.85.

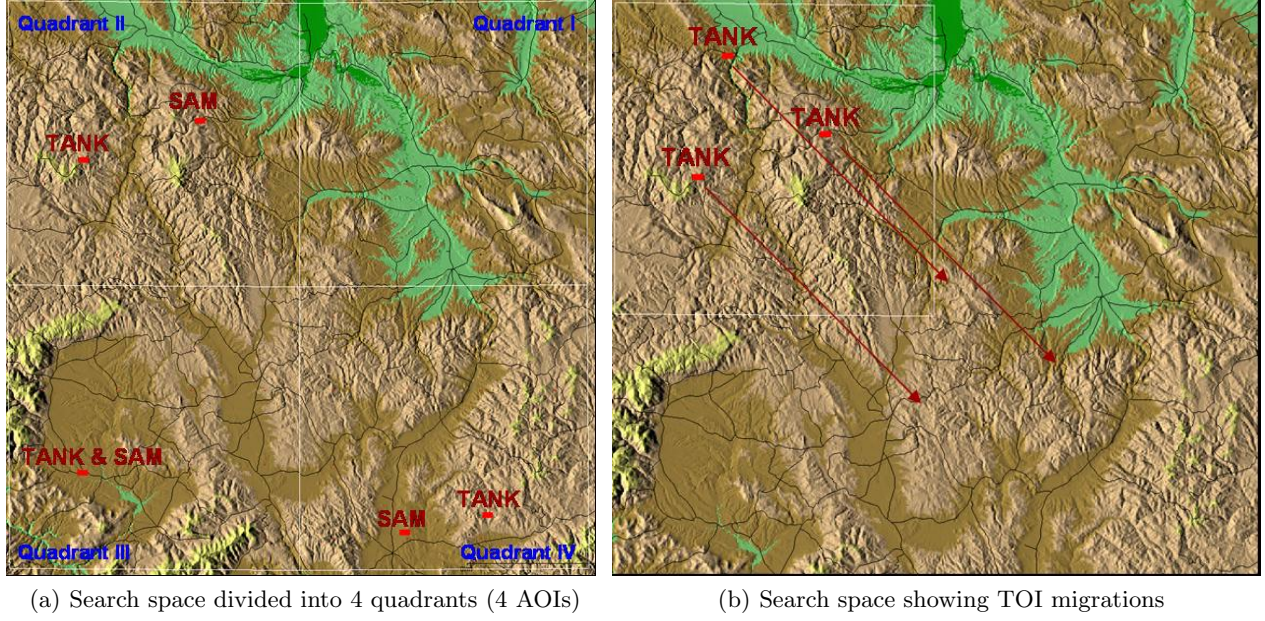


Figure 3. SLAMEMTM screen shots for cases I and II

5.1. Test Scenario I: TOIs Divided into Quadrants

In the following three test scenarios, six ground vehicles existed in isolated sub-regions of the search space. One tank and one SAM battery resided in the northwest corner (Quadrant II in Figure 3(a)), and identical pairs of vehicles resided in the southwest and southeast corners of the space (Quadrants III and IV in Figure 3(a)). There were no vehicles in the northeast corner of the search space (Quadrant I). Over the 3-hour simulation, each pair of vehicles executed move-stop and stop-move sequences that were (partially) temporally staggered from one sub-region to the next. The airborne assets in these tests carried only radar sensors, which detect and track *moving* targets. Thus, tracks were usually lost while the vehicles were stopped. The pairs of vehicles did not necessarily move together, so ambiguities in vehicle classification could arise if the vehicle paths intersected. All six of these ground vehicles were classified as TOIs, and therefore given higher tracking priority by weighting their entropy higher in the objective function.

5.1.1. Scenario I-a: No WAS

The first test case used no intelligent wide area surveillance technique. The objective function was purely track-based (J_T only), and if at any point in time a particular sensor had no track to look at, the asset manager (AM) chose the first candidate task in the list of possible tasks to be executed by the sensor. While vehicles were in track and moving, the track-based tasking did a reasonable job of maintaining track on them. However, when the vehicles transitioned from stopped to moving, the track-based tasks did not detect them very quickly, as most tasks were focused on the moving vehicles. As expected, fewer tracks were maintained in this test case. Tracking results appear in blue in Figure 4(a). Over the 3-hour simulation, an average of 2.68 TOIs were kept in track and correctly classified (displayed as the single blue diamond on the far right of the graph).

5.1.2. Scenario I-b: WAS

The second test case used the entropy-based WAS technique described in the paper ($J_T + J_W$ objective function), implemented on a discretization of the search space. The search space (approx. 200km x 200km square) was

divided into 64 equal gridblocks. A target mass vector was assumed for each gridblock: our a priori assumption $v = [0.0525, 0.0525, 0.0]$ represented a target mass distribution of 0.0525 tanks, 0.0525 SAMs, and 0.0 civilians in each gridblock. Note that this assumed target mass corresponds to a perception that 3.36 tanks and 3.36 SAMs existed in the entire search space: $(0.0525 \text{ vehicles per gridblock}) * (64 \text{ gridblocks}) = 3.36 \text{ vehicles}$. If this a priori target mass assumption were less accurate, performance of the tasking algorithms may suffer slightly. We observed more intelligently placed sensor footprints in this test case than in Case I-a, and a subsequently higher number of targets discovered and maintained in track. Tracking results appear in pink in Figure 4(a). Over the 3-hour simulation, an average of 3.30 TOIs were kept in track and correctly classified (displayed as the single pink square on the far right of the graph).

5.1.3. Scenario I-c: WAS with a priori Knowledge

The final test case is a more sophisticated version of the second. In this case, the tasking algorithm was provided with a more precise estimate of the target mass distribution. Indeed, no target mass was assumed for the northeast quadrant of the search space (reflecting a priori knowledge that no actual targets exist there). The remaining regions had uniformly distributed target mass vectors $v = [0.07, 0.07, 0.0]$ representing an assumption that 3.36 Tanks and SAMs existed in the search space. Tracking results for this case were slightly better than in Case II, since the sensor could do a more refined search for the targets. Tracking results appear in yellow in Figure 2. Over the 3-hour simulation, an average of 3.90 TOIs were kept in track and correctly classified (displayed as the single yellow triangle on the far right of the graph).

Coverage maps for Cases Ia-Ic appear as Figures 5(a)–5(c). The color scale for these maps is displayed as Figure 5(d). Figure 5(a) shows sensor coverage during Case I-a, when no intelligent WAS method was employed. Coverage of the search space is clearly non-uniform, but in no systematic way. Figure 5(b) shows sensor coverage during Case I-b. Uniform coverage of the search space was achieved by presetting the AM with a uniform target mass map. Figure 5(c) shows sensor coverage during Case I-c, when the AM was loaded with a more precise target mass map reflecting the absence of targets in Quadrant I. Figure 5(c) displays significantly lower sensor coverage of Quadrant I because of the a priori knowledge that no target existed there. In all three figures, coverage of Quadrant III is less consistent; this is an artifact of the positions of the airborne radar assets. The fixed flight paths of the planes often prevented them from viewing Quadrant III.

5.2. Test Scenario II: Staggered TOI Migration

In the following two test scenarios, three tanks move from the northwest corner of the search space (Quadrant II in Figure 6) to the southeast corner of the space (Quadrant IV). These movements are temporally staggered: the first tank begins moving immediately, the second tank begins moving 3 hours later, and the final tank begins moving 6 hours later. Since the airborne assets in these tests carried only MTI sensors, vehicles were not detected until they began to move. The vehicles continued to move around within Quadrant IV after they arrived there, rendering them trackable if the sensors had clear line-of-sight to them. The simulation runs for 8 hours.

5.2.1. Scenario II-a: No WAS

As in Case I-a, this test case used no intelligent wide area surveillance (WAS) technique. The objective function was purely track-based (J_T only). While tanks were in track and moving, the track-based tasking did a reasonable job of maintaining track on them. However, when the tanks first began moving, especially the second and third tanks, the track-based tasks did not detect them very quickly. Indeed, the first tank moved out of Quadrant II before the second tank began moving; therefore, the AM was focusing nearly all the sensor resources on the discovered (moving) tank in Quadrant IV. It is by chance that sensor tasks cover Quadrant II again and eventually detect the second and third tanks. Tracking results appear in blue in Figure 4(b). Over the 8-hour simulation, an average of 0.88 TOIs were kept in track and correctly classified (displayed as the single blue diamond on the far right of the graph).

5.2.2. Scenario II-b: WAS with a priori Knowledge

This test case implemented the entropy-based WAS technique ($J_T + J_I$ objective function) on a discretized area of interest (AOI). The WAS region was chosen to be Quadrant II only, reflecting a priori knowledge by the AM that all the tanks existed in that region. The WAS region was divided into 16 equal grid-blocks and the target mass vector $v = [0.0, 0.2, 0.0]$ was assumed for each grid-block, reflecting an assumption that 3.2 tanks (and no other vehicles) existed in the AOI ($16 * 0.2 = 3.2$). Instead of following the first tank out of Quadrant II and rarely looking back, as in Case IV, the AM continues to task sensors to cover Quadrant II (the AOI) even while other vehicles are in track elsewhere. This persistent coverage of Quadrant II allows the AM to discover the other tanks much sooner after they begin moving. Tracking results appear in pink in Figure 4(b). Over the 3-hour simulation, an average of 1.46 TOIs were kept in track and correctly classified (displayed as the single pink square on the far right of the graph).

Coverage maps for Cases II-a and II-b appear below as Figures 6(a) and 6(b). Figure 6(a) shows sensor coverage during Case II-a, when no intelligent WAS method was employed. Coverage is heavier in Quadrant IV because the sensors followed the first tank to that region and continued to view that region almost exclusively. The tanks remaining in Quadrant II were stationary for most of the simulation, remained undiscovered to the AM, and gave the AM no reason to task sensors there. The tank(s) in Quadrant IV kept moving until the end of the simulation and therefore remained trackable. Figure 6(b) shows sensor coverage during Case II-b. Coverage is heavier in Quadrant II, which is good, considering that the majority of the tanks spent the majority of their time there. The WAS region, laden with target mass, motivated persistent coverage of Quadrant II, even though the tanks that existed there were often sitting still and therefore undetectable.




Figure 4. TOIs correctly classified over time

6. SUMMARY

The dynamic tasking algorithm described in this paper evaluates tasks on their ability to improve knowledge of existing tracks as well to acquire new tracks. As can be seen from these test scenarios, neglecting WAS leads to significantly poorer tracking performance. We observed that schedules designed to minimize entropy of existing tracks tend to focus sensor resources on small geographical areas. This is undesirable for tracking problems involving extended operating conditions such as multiple stop-move-stop cycles and long mission durations. If, during the course of a mission, track of a TOI is lost, the focusing of resources only on existing tracks will significantly reduce the probability of reacquiring the TOI at a later time. By requiring the objective function to also have a track initiation component J_W , we are forcing the algorithm to strike a balance between track-focused sensor coverage and WAS coverage designed to acquire (and reacquire) tracks.

ACKNOWLEDGMENTS

This research was sponsored in part by the Air Force Research Laboratory under Contract F33615-03-C-1454 and by the Navy under Contract N68335-03-C-0241. The authors acknowledge Chris Heithecker for his support.

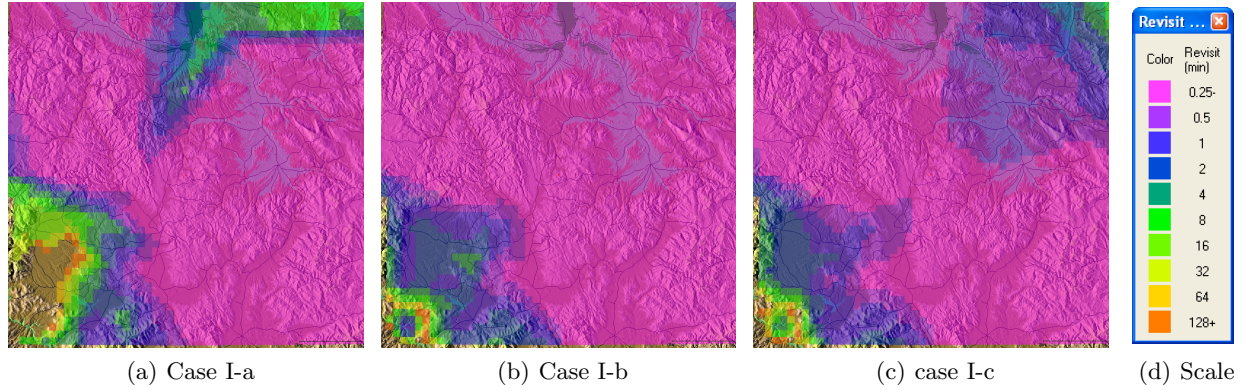


Figure 5. Case I: GMTI coverage over 3-hour scenario

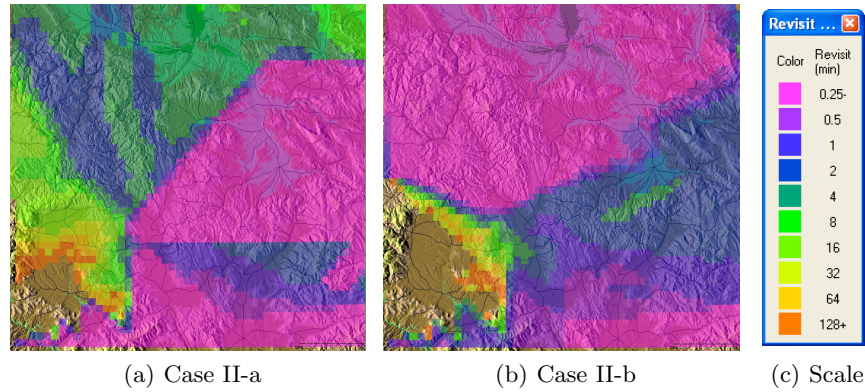


Figure 6. Case II: GMTI coverage over 3-hour scenario

REFERENCES

1. K. Kastella, "Discrimination gain to optimize detection and classification," *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans* **27**, pp. 112–116, January 1997.
2. G. McIntyre and K. Hintz, "An information theoretic approach to sensor scheduling," in *Proceedings of the SPIE International Symposium on Aerospace/Defense Sensing and Control*, **2755**, pp. 304–312, 1996.
3. G. McIntyre and K. Hintz, "Sensor management simulation and comparative study," in *Proceedings of the SPIE International Symposium on Aerospace/Defense Sensing and Control*, **3068**, p. 250260, 1997.
4. C. Shannon, "A mathematical theory of communication," *Bell System Technical Journal* **27**, pp. 379–423, July 1948.
5. N. Xiong and P. Svensson, "Multi-sensor management for information fusion: issues and approaches," *Information Fusion* **3**, pp. 163–186, 2002.
6. A. Steinberg, C. Bowman, and F. White Jr., "Revisions to the jdl data fusion model," in *Proc. 3rd NATO/IRIS Conf.*, (Quebec City, Canada), 1998.
7. S. Musick and R. Malhotra, "Chasing the elusive sensor manager," in *Proceedings of the National Aerospace Electronics Conference (NAECON)*, (Dayton, OH), May 1994.
8. K. Hintz and E. McVey, "Multi-process constrained estimation," *IEEE Transaction on Systems, Man, and Cybernetics* **21**, pp. 434–442, Jan/Feb 1991.
9. M. Meloon and K. Ezal, *An Entropy-Based Approach to Opportunistic Sensor Management*, March 2003.
10. M. Meloon, R. Cagley, and P. Stieber, *An Entropy-Based Approach to Opportunistic Sensor Resource Management*, July 2004.